

AlterNIC
The Network Information Center.
<http://www.alternic.org/>
Because everything else just ends in .COM

Send Comments: <mailto:hostmaster@alternic.org>

RFC Search Engine.
<http://www.alternic.org/rfcs/search/>

rfc2267.txt

Network Working Group
Request for Comments: 2267
Category: Informational

P. Ferguson
Cisco Systems, Inc.
D. Senie
BlazeNet, Inc.
January 1998

Network Ingress Filtering:
Defeating Denial of Service Attacks which employ
IP Source Address Spoofing

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1998). All Rights Reserved.

Abstract

Recent occurrences of various Denial of Service (DoS) attacks which have employed forged source addresses have proven to be a troublesome issue for Internet Service Providers and the Internet community overall. This paper discusses a simple, effective, and straightforward method for using ingress traffic filtering to prohibit DoS attacks which use forged IP addresses to be propagated from 'behind' an Internet Service Provider's (ISP) aggregation point.

Table of Contents

1. Introduction	2
2. Background	2
3. Restricting forged traffic	5
4. Further capabilities for networking equipment.	6
5. Liabilities.	6
6. Summary.	7
7. Security Considerations.	7
8. Acknowledgments	8
9. References	8

10. Authors' Addresses	9
11. Full Copyright Statement	10

1. Introduction

A resurgence of Denial of Service Attacks [1] aimed at various targets in the Internet have produced new challenges within the Internet Service Provider (ISP) and network security communities to find new and innovative methods to mitigate these types of attacks. The difficulties in reaching this goal are numerous; some simple tools already exist to limit the effectiveness and scope of these attacks, but they have not been widely implemented.

This method of attack has been known for some time. Defending against it, however, has been an ongoing concern. Bill Cheswick is quoted in [2] as saying that he pulled a chapter from his book, "Firewalls and Internet Security" [3], at the last minute because there was no way for an administrator of the system under attack to effectively defend the system. By mentioning the method, he was concerned about encouraging it's use.

While the filtering method discussed in this document does absolutely nothing to protect against flooding attacks which originate from valid prefixes (IP addresses), it will prohibit an attacker within the originating network from launching an attack of this nature using forged source addresses that do not conform to ingress filtering rules. All providers of Internet connectivity are urged to implement filtering described in this document to prohibit attackers from using forged source addresses which do not reside within a range of legitimately advertised prefixes. In other words, if an ISP is aggregating routing announcements for multiple downstream networks, strict traffic filtering should be used to prohibit traffic which claims to have originated from outside of these aggregated announcements.

An additional benefit of implementing this type of filtering is that it enables the originator to be easily traced to it's true source, since the attacker would have to use a valid, and legitimately reachable, source address.

2. Background

A simplified diagram of the TCP SYN flooding problem is depicted below:

```

                                     9.0.0.0/8
host <----- router <--- Internet <----- router <-- attacker

```

TCP/SYN

<-----
Source: 192.168.0.4/32

Ferguson & Senie

Informational

[Page 2]

RFC 2267

Network Ingress Filtering

January 1998

SYN/ACK

no route

TCP/SYN

<-----
Source: 10.0.0.13/32

SYN/ACK

no route

TCP/SYN

<-----
Source: 172.16.0.2/32

SYN/ACK

no route

[etc.]

Assume:

- o The "host" is the targeted machine.
- o The attacker resides within the "valid" prefix, 9.0.0.0/8.
- o The attacker launches the attack using randomly changing source addresses; in this example, the source addresses are depicted as from within [4], which are not generally present in the global Internet routing tables, and therefore, unreachable. However, any unreachable prefix could be used to perpetrate this attack method.

Also worthy of mention is a case wherein the source address is forged to appear to have originated from within another legitimate network which appears in the global routing table(s). For example, an attacker using a valid network address could wreak havoc by making the attack appear to come from an organization which did not, in fact, originate the attack and was completely innocent. In such cases, the administrator of a system under attack may be inclined to filter all traffic coming from the apparent attack source. Adding such a filter would then result in a denial of service to legitimate, non-hostile end-systems. In this case, the administrator of the system under attack unwittingly becomes an accomplice of the attacker.

Further complicating matters, TCP SYN flood attacks will result in SYN-ACK packets being sent to one or many hosts which have no involvement in the attack, but which become secondary victims. This allows the attacker to abuse two or more systems at once.

Similar attacks have been attempted using UDP and ICMP flooding. The former attack (UDP flooding) uses forged packets to try and connect the chargen UDP service to the echo UDP service at another site. Systems administrators should NEVER allow UDP packets destined for system diagnostic ports from outside of their administrative domain to reach their systems. The latter attack (ICMP flooding), uses an insidious feature in IP subnet broadcast replication mechanics. This attack relies on a router serving a large multi-access broadcast network to frame an IP broadcast address (such as one destined for 10.255.255.255) into a Layer 2 broadcast frame (for ethernet, FF:FF:FF:FF:FF:FF). Ethernet NIC hardware (MAC-layer hardware, specifically) will only listen to a select number of addresses in normal operation. The one MAC address that all devices share in common in normal operation is the media broadcast, or FF:FF:FF:FF:FF:FF. In this case, a device will take the packet and send an interrupt for processing. Thus, a flood of these broadcast frames will consume all available resources on an end-system [9]. It is perhaps prudent that system administrators should consider ensuring that their border routers do not allow directed broadcast packets to be forwarded through their routers as a default.

When an TCP SYN attack is launched using unreachable source address, the target host attempts to reserve resources waiting for a response. The attacker repeatedly changes the bogus source address on each new packet sent, thus exhausting additional host resources.


Alternatively, if the attacker uses someone else's valid host address as the source address, the system under attack will send a large number of SYN/ACK packets to what it believes is the originator of the connection establishment sequence. In this fashion, the attacker does damage to two systems: the destination target system, as well as the system which is actually using the spoofed address in the global routing system.

The result of both attack methods is extremely degraded performance, or worse, a system crash.

In response to this threat, most operating system vendors have modified their software to allow the targeted servers to sustain attacks with very high connection attempt rates. This is a welcome and necessary part of the solution to the problem. Ingress filtering will take time to be implemented pervasively and be fully effective, but the extensions to the operating systems can be implemented quickly. This combination should prove effective against source address spoofing. See [1] for vendor and platform software upgrade information.

3. Restricting forged traffic

The problems encountered with this type of attack are numerous, and involve shortcomings in host software implementations, routing methodologies, and the TCP/IP protocols themselves. However, by restricting transit traffic which originates from a downstream network to known, and intentionally advertised, prefix(es), the problem of source address spoofing can be virtually eliminated in this attack scenario.



In the example above, the attacker resides within 9.0.0.0/8, which is provided Internet connectivity by ISP D. An input traffic filter on the ingress (input) link of "router 2", which provides connectivity to the attacker's network, restricts traffic to allow only traffic originating from source addresses within the 9.0.0.0/8 prefix, and prohibits an attacker from using "invalid" source addresses which reside outside of this prefix range.

In other words, the ingress filter on "router 2" above would check:

```

IF    packet's source address from within 9.0.0.0/8
THEN  forward as appropriate

IF    packet's source address is anything else
THEN  deny packet
  
```

Network administrators should log information on packets which are dropped. This then provides a basis for monitoring any suspicious activity.

4. Further possible capabilities for networking equipment

Additional functions should be considered for future platform implementations. The following one is worth noting:

- o Implementation of automatic filtering on remote access servers. In most cases, a user dialing into an access server is an individual user on a single PC. The ONLY valid source IP address for packets originating from that PC is the one assigned by the ISP (whether statically or dynamically assigned). The remote access server could check every packet on ingress to ensure the user is not spoofing the source address on the packets which he is originating. Obviously, provisions also need to be made for cases where the customer legitimately is attaching a net or subnet via a remote router, but this could certainly be implemented as an optional parameter. We have received reports that some vendors and some ISPs are already starting to implement this capability.

We considered suggesting routers also validate the source IP address of the sender as suggested in [8], but that methodology will not operate well in the real networks out there today. The method suggested is to look up source addresses to see that the return path to that address would flow out the same interface as the packet arrived upon. With the number of asymmetric routes in the Internet, this would clearly be problematic.

5. Liabilities

Filtering of this nature has the potential to break some types of "special" services. It is in the best interest of the ISP offering these types of special services, however, to consider alternate methods of implementing these services to avoid being affected by ingress traffic filtering.

Mobile IP, as defined in [6], is specifically affected by ingress traffic filtering. As specified, traffic to the mobile node is tunneled, but traffic from the mobile node is not tunneled. This results in packets from the mobile node(s) which have source addresses that do not match with the network where the station is attached. The Mobile IP Working Group is addressing this problem by specifying "reverse tunnels" in [7]. This work in progress provides a method for the data transmitted from the mobile node to be tunneled to the home agent before transmission to the Internet. There are additional benefits to the reverse tunneling scheme, including better handling of multicast traffic. Those implementing mobile IP systems are encouraged to implement this method of reverse tunneling.

As mentioned previously, while ingress traffic filtering drastically

reduces the success of source address spoofing, it does not preclude an attacker using a forged source address of another host within the permitted prefix filter range. It does, however, ensure that when an attack of this nature does indeed occur, a network administrator can be sure that the attack is actually originating from within the known prefixes that are being advertised. This simplifies tracking down the culprit, and at worst, the administrator can block a range of source addresses until the problem is resolved.

If ingress filtering is used in an environment where DHCP or BOOTP is used, the network administrator would be well advised to ensure that packets with a source address of 0.0.0.0 and a destination of 255.255.255.255 are allowed to reach the relay agent in routers when appropriate. The scope of directed broadcast replication should be controlled, however, and not arbitrarily forwarded.

6. Summary

Ingress traffic filtering at the periphery of Internet connected networks will reduce the effectiveness of source address spoofing denial of service attacks. Network service providers and administrators have already begun implementing this type of filtering on periphery routers, and it is recommended that all service providers do so as soon as possible. In addition to aiding the Internet community as a whole to defeat this attack method, it can also assist service providers in locating the source of the attack if service providers can categorically demonstrate that their network already has ingress filtering in place on customer links.

Corporate network administrators should implement filtering to ensure their corporate networks are not the source of such problems. Indeed, filtering could be used within an organization to ensure users do not cause problems by improperly attaching systems to the wrong networks. The filtering could also, in practice, block a disgruntled employee from anonymous attacks.

It is the responsibility of all network administrators to ensure they do not become the unwitting source of an attack of this nature.

7. Security Considerations

The primary intent of this document is to inherently increase security practices and awareness for the Internet community as a whole; as more Internet Providers and corporate network administrators implement ingress filtering, the opportunity for an attacker to use forged source addresses as an attack methodology will significantly lessen. Tracking the source of an attack is simplified

when the source is more likely to be "valid." By reducing the number and frequency of attacks in the Internet as a whole, there will be more resources for tracking the attacks which ultimately do occur.

8. Acknowledgments

The North American Network Operators Group (NANOG) [5] group as a whole deserves special credit for openly discussing these issues and actively seeking possible solutions. Also, thanks to Justin Newton [Priori Networks] and Steve Bielagus [OpenROUTE Networks, Inc.] for their comments and contributions.

9. References

- [1] CERT Advisory CA-96.21; TCP SYN Flooding and IP Spoofing Attacks; September 24, 1996.
- [2] B. Ziegler, "Hacker Tangles Panix Web Site", Wall Street Journal, 12 September 1996.
- [3] "Firewalls and Internet Security: Repelling the Wily Hacker"; William R. Cheswick and Steven M. Bellovin, Addison-Wesley Publishing Company, 1994; ISBN 0-201-63357-4.
- [4] Rekhter, Y., Moskowitz, R., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", RFC 1918, February 1996.
- [5] The North American Network Operators Group;
<http://www.nanog.org>.
- [6] Perkins, C., "IP Mobility Support", RFC 2002, October 1996.
- [7] Montenegro, G., "Reverse Tunneling Mobile IP",
Work in Progress.
- [8] Baker, F., "Requirements for IP Version 4 Routers", RFC 1812,
June 1995.
- [9] Thanks to: Craig Huegen;
See: <http://www.quadrunner.com/~chuegen/smurf.txt>.

10. Authors' Addresses

Paul Ferguson
Cisco Systems, Inc.
400 Herndon Parkway
Herndon, VA USA 20170

E-Mail: ferguson@cisco.com

Daniel Senie
BlazeNet, Inc.
4 Mechanic Street
Natick, MA USA 01760

E-Mail: dts@senie.com

11. Full Copyright Statement

Copyright (C) The Internet Society (1998). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this

document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.